arXiv:1605.02637v1 [math.NT] 9 May 2016

A DATABASE OF HILBERT MODULAR FORMS

STEVE DONNELLY AND JOHN VOIGHT

Abstract. We describe the computation of tables of Hilbert modular forms of parallel weight 2 over totally real fields.

Tables of classical modular forms have been computed for many decades now, starting with a table of rational eigenforms (corresponding to modular elliptic curves) appearing in the proceedings of the Antwerp conference in 1972 [1]. Large tables of classical modular forms now exist, as computed by Cremona [6, 7] for rational eigenforms, Stein [16, 17], and more recently appearing in the *L-functions and Modular Forms Database* (LMFDB) [14], computed by Stephan Ehlen and Fredrik Strömberg. In this article, we continue this tradition of tabulating automorphic forms on Shimura varieties by computing extensive tables of Hilbert modular cusp forms of parallel weight 2 over totally real fields. This work is in a sense complementary to more recent work [2] on tables of (modular) elliptic curves over $\mathbb{Q}(\sqrt{5})$. Our data is available at the second author's webpage [18] and is searchable online at the LMFDB [14].

In Section 1, we give an overview of the basic algorithms employed. In Section 2, we describe in more detail the implementation in the computer algebra system Magma [3], providing some documentation and detail about computations in practice and a few auxiliary algorithms employed. In Section 3, we gather the data which was computed and discuss the extent of the computation. Finally, in Section 4 we some interesting aspects of the data in detail.

The authors would like to thank Lassina Dembélé and Matthew Greenberg for useful discussions. Thanks also go to John Cremona and Aurel Page for adding Galois conjugate Hilbert modular form data to the LMFDB. The second author was supported by an NSF Grant (DMS-0901971) during the time that these computations were first undertaken, and by an NSF CAREER Award (DMS-1151047) while work was completed.

1. The algorithm: an overview

In this section, we give a basic overview of the algorithm we employ. For a motivated introduction, see the expository work by Dembélé–Voight [10]; for a more basic reference on Hilbert modular forms, see e.g. Freitag [11].

Throughout, let F be a totally real field of degree $n = [F : \mathbb{Q}]$ with ring of integers $\mathbb{Z}_F$ and let $\mathfrak{N} \subseteq \mathbb{Z}_F$ be an ideal. Let $S_k(\mathfrak{N})$ denote the space of Hilbert cusp forms of level $\mathfrak{N}$ and parallel weight $k \in \mathbb{Z}_{\geq 2}$ for F, and let $S_k^{\mathrm{new}}(\mathfrak{N})$ be the new subspace.

The algorithm we use combines methods of Dembélé [8] and Dembélé–Donnelly [9], along with methods of Greenberg–Voight [12] and Voight [20].

Theorem 1.1. *There exists an (explicit) algorithm that, on input a totally real field F and an ideal $\mathfrak{N} \subseteq \mathbb{Z}_F$, computes as output the systems of eigenvalues for the Hecke operators $T_{\mathfrak{p}}$ with $\mathfrak{p} \nmid \mathfrak{N}$ and the Atkin-Lehner involutions $W_{\mathfrak{p}^e}$ with $\mathfrak{p}^e \parallel \mathfrak{N}$ that occur in the space $S_k^{\mathrm{new}}(\mathfrak{N})$.*

In other words, there exists an explicit finite procedure which takes as input the totally real field F, an ideal $\mathfrak{N} \subseteq \mathbb{Z}_F$ (encoded in bits in the usual way [5]), and an even integer $k \geq 2$, and outputs a finite set of number fields $E_f \subseteq \overline{\mathbb{Q}}$ and sequences $(a_f(\mathfrak{p}))_{\mathfrak{p}}$ encoding the Hecke eigenvalues for each eigenform constituent f in $S_k(\mathfrak{N})$ and $a_f(\mathfrak{p}) \in E_f$.

In this article, we consider the first and essential case $k = 2$; we believe it would be an interesting project to extend these efforts to include tabulation of Hilbert modular forms of higher weight as well.

The general strategy of the algorithm in Theorem 1.1 is as follows. Let B be the quaternion algebra over F which is unramified at all finite places and ramified at either all or all but one real place according as $n = [F : \mathbb{Q}]$ is even or odd. Then by the Eichler–Shimizu–Jacquet–Langlands correspondence, there exists an isomorphism of Hecke modules

$$S_2^B(\mathfrak{N}) \xrightarrow{\sim} S_2(\mathfrak{N})$$

where $S_2^B(\mathfrak{N})$ denotes the space of quaternionic modular forms over B of weight k and level $\mathfrak{N}$. The description of the space $S_2^B(\mathfrak{N})$ varies accordingly as n is even or odd. Let $\mathcal{O} \subseteq B$ be a maximal order.

If n is even, we use Brandt matrices, following the method of Dembélé [8] and Dembéle–Donnelly [9]. The Hecke module $S_2^B(\mathfrak{N})$ is computed via an explicit action on maps from the set of right ideal classes in $\mathcal{O}$ to $\mathbb{P}^1(\mathbb{Z}_F/\mathfrak{N})$. We refer to this as the definite method, since the algebra B in this case is definite.

If instead n is odd, the algorithm in Theorem 1.1 is supplied by Greenberg–Voight [12] and Voight [20]. For simplicity of exposition, assume F has strict class number 1; the method works for arbitrary class number. Let $\mathcal{O}_1^\times = \{\gamma \in \mathcal{O} : \mathrm{nrd}(\gamma) = 1\}$ denote the group of units of $\mathcal{O}$ with reduced norm 1. Suppose that the unique real split place of B corresponds to the embedding $\iota_\infty : B \hookrightarrow \mathrm{M}_2(\mathbb{R})$ and let $\Gamma(1) = \iota_\infty(\mathcal{O}_1^\times)/\{\pm 1\} \subseteq \mathrm{PSL}_2(\mathbb{R})$. Let $\mathcal{O}_0(\mathfrak{N})$ be an Eichler order of level $\mathfrak{N}$ and let $\Gamma_0(\mathfrak{N}) = \iota_\infty(\mathcal{O}_0(\mathfrak{N})_1^\times)/\{\pm 1\}$. Then by the theorem of Eichler-Shimura, and Shapiro's lemma, we have a further isomorphism of Hecke modules

$$S_2^B(\mathfrak{N}) \xrightarrow{\sim} H^1\big(\Gamma(1), \mathrm{Coind}_{\Gamma_0(\mathfrak{N})}^{\Gamma(1)} \mathbb{C}\big)^+$$

where $+$ denotes the $+1$-eigenspace for complex conjugation. The Hecke module $S_2^B(\mathfrak{N})$ is then computed via this isomorphism as group cohomology, using an explicit presentation for the group $\Gamma(1)$ based on the computation of a fundamental domain [19]. We call this the indefinite method.

Remark 1.2. More generally, one can use a quaternion algebra B over F of discriminant $\mathfrak{D}$, ramified at all or all but one real place, and $\mathfrak{M}$ an ideal coprime to $\mathfrak{D}$; then with $\mathfrak{N} = \mathfrak{D}\mathfrak{M}$, the Jacquet-Langlands correspondence yields an isomorphism

$$S_2^B(\mathfrak{M}) \xrightarrow{\sim} S_2(\mathfrak{N})^{\mathfrak{D}\text{-new}}$$

where $S_2(\mathfrak{N})^{\mathfrak{D}\text{-new}} \subseteq S_2(\mathfrak{N})$ denotes the space of Hilbert cusp forms which are new at all primes dividing $\mathfrak{D}$. Since $\mathfrak{D}$ is squarefree, these methods overlap when there is a prime $\mathfrak{p}$ which exactly divides the level $\mathfrak{N}$. We use this to verify our calculations by computing with both the definite and indefinite method when it applies; for more, see Section 3.

2. Computing systems of Hecke eigenvalues

In this section, we discuss the Magma [3] implementation of the above algorithms. These methods saw a first implementation over $\mathbb{Q}$ by David Kohel, and they rely heavily on code for quaternion algebras over number fields, due to Donnelly, Markus Kirschmer, Nicole Sutherland, and Voight (documented in work of Kirschmer–Voight [13]).

Our algorithm has 6 main steps:

1. Precomputation
2. Computing a basis
3. Hecke computation via reduction
4. Decomposition into newforms

We refer to Dembélé–Voight [10] for an explicit mathematical description, and here we provide some additional comments about the implementation.

Step 1: Precomputation. The first step of the algorithm involves precomputation steps that only depend on the field F: that is to say, for all levels $\mathfrak{N}$ with $\mathfrak{p} \parallel \mathfrak{N}$, this step need only be performed once. For the definite method, the precomputation step involves determining a set of representatives for the right ideal classes in $\mathcal{O}$ up to isomorphism.

- We precompute the theta series of each left order $\mathcal{O}$ in the set of ideal representatives, recording
$$\#\{\alpha \in \mathcal{O}_L(I) : \mathrm{Tr}_{F/\mathbb{Q}}\,\mathrm{nrd}(\alpha) = n\}$$
for a small set of values $n \in \mathbb{Z}_{\geq 0}$. This reduces the number of isomorphism tests of ideals when computing Hecke operators below, but at the marginal additional cost of requiring the computation of left orders and representation numbers.
- As computing inverses and colon ideals can also be time consuming, we compute for each representative right ideal I a small element $b \in I$ and a left ideal I' of prime power norm such that $\mathcal{O}_L(I)b = II'$: then $JI' = J(I^{-1})b$, and α generates JI' if and only if α/b generates JI^{-1}.

For the indefinite method, we compute a fundamental domain for the Shimura curve associated to the maximal order, together with a presentation for the group $\Gamma_0(1)$ [19].

Step 2: Computing a basis. To compute a basis of forms, we need fast algorithms for working with $\mathbb{P}^1(\mathbb{Z}_F/\mathfrak{N})$, and our efficient implementation treats the case where $\mathfrak{N}$ is prime and composite differently: in the latter we index elements $[a : b] \in \mathbb{P}^1(\mathbb{Z}_F/\mathfrak{N})$ according to divisor $\mathfrak{a} = (a) \mid \mathfrak{N}$. With these representatives in hand, in the definite case we list orbits of $\mathbb{P}^1(\mathbb{Z}_F/\mathfrak{N})$ under the action of the stabilizer groups and in the indefinite case we record the permutation action in the coinduced module. (If we were working in higher weight, and a nontrivial representation V, we would put a copy of V on each element of $\mathbb{P}^1(\mathbb{Z}_F/\mathfrak{N})$.)

For sanity, we also verify the dimension of the space using a dimension formula (involving class numbers and embedding numbers).

Step 3: Hecke computation via reduction. For the computation of Hecke operators, we use reduction on $\mathbb{P}^1(\mathbb{Z}_F/\mathfrak{N})$ in the definite case and reduction theory in the fundamental domain in the indefinite case.

Step 4: Decomposition into newforms. With Hecke matrices in hand, we now outline the procedures used in our implementation to decompose a newspace as a direct sum of new simultaneous eigenspaces. Let $S = S_2(\mathfrak{N})$. We use the following properties.

(1) The Hecke operators $T_{\mathfrak{p}}$ for $\mathfrak{p}$ coprime to $\mathfrak{N}$ all commute.
(2) The Hecke eigenspaces in S have distinct systems of eigenvalues for the operators $T_{\mathfrak{p}}$ with $\mathfrak{p}$ coprime to $\mathfrak{N}$ ("multiplicity one").
(3) We have $\dim S^{\mathfrak{p}\text{-old}} = 2\dim S_2(\mathfrak{N}/\mathfrak{p}) - \dim S_2(\mathfrak{N}/\mathfrak{p}^2)$, with the last term omitted if $\mathfrak{p}^2$ does not divide $\mathfrak{N}$.

These properties are standard: for (2), see for instance work of Shemanske and Walling [15, Theorem 3.6], and (3) follows from standard facts about degeneracy maps (the proof is a straightforward generalization of the classical proof (where $F = \mathbb{Q}$), and is omitted).

Using (1) and standard linear algebra techniques, we decompose the space into Hecke constituents corresponding to $\mathbb{Q}$-irreducible subspaces (a direct sum of eigenspaces). As harmless as it may seem, this step is often the bottleneck in computation.

Finally, we show how to decompose a space into old and new subspaces. For a prime $\mathfrak{p} \nmid \mathfrak{N}$, the $\mathfrak{p}$-old subspace $S^{\mathfrak{p}\text{-old}}$ of $S = S_2(\mathfrak{N})$ is the span of the images of all degeneracy maps $S_2(\mathfrak{N}/\mathfrak{p}) \to S_2(\mathfrak{N})$.

The old subspace S^{old} is the sum of the spaces $S^{\mathfrak{p}\text{-old}}$ for primes $\mathfrak{p} \mid \mathfrak{N}$. The newspaces S^{new} and $S^{\mathfrak{p}\text{-new}}$ are the unique Hecke-invariant complements of S^{old} and $S^{\mathfrak{p}\text{-old}}$, respectively.

Now suppose we are given two Hecke-modules S and S_0 (e.g. $S_0 = S_2(\mathfrak{N}/\mathfrak{p})$ and we wish to compute the Hecke-submodule S^{old} of S consisting of all the irreducible Hecke-submodules (i.e. simultaneous eigenforms) of S that arise from S_0 via degeneracy maps. In addition, suppose we know a priori the dimension of S^{old}, as in Property (3). For a prime $\mathfrak{p}$ coprime to $\mathfrak{N}$, the space S^{old} is contained in the kernel of the characteristic polynomial of $T_\mathfrak{p} \mid S_0$ evaluated at $T_\mathfrak{p} \mid S$. Since we know the dimension, we may compute S^{old} in this manner using finitely many primes $\mathfrak{p}$, stopping when the intersection has the right dimension. One could obtain the complement S^{new} by computing the Hecke-invariant complement, requiring extra module-splitting effort. However, for present purposes, we do not need to know S^{new} as a particular subspace of S; we merely need the Hecke action on S^{new}, which can be obtained using the quotient $S/S^{\text{old}} \cong S^{\text{new}}$.

Alternatively, we can compute the degeneracy operators themselves, and then compute the orthogonal complement.

Now suppose that S is a new subspace of dimension d. By (1), it contains d simultaneous eigenvectors having distinct systems of eigenvalues, which we wish to determine. The Hecke algebra $\mathbb{T} = \mathbb{Z}[T_\mathfrak{n}]_\mathfrak{n}$ is a commutative $\mathbb{Z}$-algebra which is generated by Hecke operators $T_\mathfrak{p}$ with $\mathfrak{p}$ coprime to $\mathfrak{N}$. Therefore $\mathbb{T} \otimes \mathbb{Q}$ is a direct product of number fields, an étale $\mathbb{Q}$-algebra generated by a single primitive element. For practical purposes, one may find a primitive element by trial and error (for instance taking simple combinations of the first few Hecke operators), since the set of non-primitive elements must be sparse. Then the simultaneous eigenspaces of S are simply the eigenvectors associated to the primitive element.

3. Data computed

We compute for the totally real fields F of degree $n \leq 6$ with discriminant $\leq 2 \cdot 10^n$. Note that if $[F : \mathbb{Q}] > 6$ then $d_F > 2 \cdot 10^n$ by the Odlyzko bounds; indeed, the totally real field of degree 7 of smallest discriminant is blank.

In this range, the discriminant and degree of the field uniquely characterizes the field, with one exceptional pair: in degree 4, there are two fields of discriminant 16448: one, with minimal polynomial $x^4 - 2x^3 - 6x^2 + 2$ has Galois group S_4 and the other $x^4 - 2x^3 - 7x^2 + 8x + 14$ has Galois group D_8. (The first multiple discriminant for cubic fields is 3969.)

This gives a total of 233676 Hilbert modular forms in the database. To save space, we did not include data for conjugate levels; for consistency across the database, these were later added to the LMFDB, increasing the overall size of the data.

The computation was performed using the Vermont Advanced Computing cluster (VACC). The odd degree fields, which used the indefinite method, took approximately 4 CPU years to compute. Much of this time is spent computing Atkin-Lehner involutions and coping with higher class number issues.

The data is hosted at the LMFDB [14] and backed up at the NECC File Exchange. The raw data takes 172 GB, compressed 78 GB.

Remark 3.1. As in Remark 1.2, in many situations we can employ either the definite or indefinite method. For example, if F has odd degree and $\mathfrak{p} \parallel \mathfrak{N}$, then one can apply the definite method with a quaternion algebra B ramified at all real places and the prime $\mathfrak{p}$; the Hecke module computed is then $S_2(\mathfrak{N})^{\mathfrak{p}\text{-new}}$. To compute a single space of forms, the definite method runs more quickly than the indefinite method, but for tabulation one must use the indefinite method if one wants to compute with forms of square level; it is also more efficient to use the indefinite method since for each quaternion algebra used one must repeat the precomputation step.

d_F	N	num	d_F	N	num
5	4999	7583	49	2059	838
8	5000	12875	81	719	343
12	5000	10565	148	499	1068
13	1999	5837	169	625	515
17	988	3387	229	256	673
	⋮			⋮	
493	4	28	1944	54	568
497	7	19	1957	27	199
		104593			40330

d_F	N	num	d_F	N	num	d_F	N	num
725	4091	2807	14641	1013	217	300125	911	177
1125	991	375	24217	835	1062	371293	961	251
	⋮			⋮			⋮	
19821	97	241	195829	83	234	1997632	127	84
		64329			15402			9022

4. Comments on data

Base change forms. We detect base change forms as follows. First, a necessary condition for f to be a base change from a smaller field E is that $\sigma(\mathfrak{N}) = \mathfrak{N}$ and $a_{\sigma(\mathfrak{p})} = a_{\mathfrak{p}}$ for all $\sigma \in \mathrm{Gal}(F/E)$. If the form fails this test (for all possible subfields E), it is not a base change form; if it passes this test, then is likely to be a base change. (One could detect more generally forms that are twists of base change forms.) To be certain, we look in the database for forms over the field E with level supported at primes dividing the discriminant $\mathfrak{d}_{F/E}$ and $\mathfrak{n} = \mathbb{Z}_E \cap \mathfrak{N}$; the precise recipe is provided by Loeffler and Weinstein for forms that are base change from $\mathbb{Q}$.

CM forms. We detect CM forms as follows. Let $\mathfrak{p}$ be a prime which is trivial in the narrow class group; then the Hecke eigenvalue $a_{\mathfrak{p}}$ is a totally real algebraic integer and the characteristic polynomial of Frobenius $T^2 - a_{\mathfrak{p}}T + N(\mathfrak{p})$ has discriminant $a_{\mathfrak{p}}^2 - 4N(\mathfrak{p})$. If the form is CM, then the Frobenius belongs to the quadratic field of this discriminant. If we find two such primes $\mathfrak{p}$ with $a_{\mathfrak{p}} \neq 0$ whose CM discriminants belong to different square classes in $F^\times/F^{\times 2}$, then the form does not have CM; if these square classes agree, then it gives strong likelihood that the form is CM with the given CM discriminant K. One can then also verify that $a_{\mathfrak{p}} = 0$ for primes $\mathfrak{p}$ inert in K to provide further evidence. To conclude that the form really is CM, we could use the machinery for Hecke characters.

Computing the L-function. Given the Hecke eigenvalues $a_{\mathfrak{p}}(f)$ of a Hilbert modular form f with trivial central character, we define its L-series via the Euler product

$$L(s,f) = \prod_{\mathfrak{p}|\mathfrak{N}} (1 - a_{\mathfrak{p}}(\mathrm{N}\mathfrak{p})^{-s})^{-1} \prod_{\mathfrak{p}\nmid\mathfrak{N}} (1 - a_{\mathfrak{p}}(\mathrm{N}\mathfrak{p})^{-s} + (\mathrm{N}\mathfrak{p})^{-2s})^{-1}.$$

The function $L(s,f)$ has analytic continuation to all of $\mathbb{C}$ and a functional equation

$$\Lambda(s,f) = (d_F^2\,\mathrm{N}(\mathfrak{N}))^{s/2}\Gamma_{\mathbb{C}}(s+1/2)^n L(s,f) = \Lambda(1-s,f).$$

Factorization of new subspaces. The largest Hecke-irreducible space in our database has dimension 286, and occurs for the field $F = \mathbb{Q}(\sqrt{296})$ and level $\mathfrak{N}$ of norm 29. It would be interesting to consider the distribution of the dimensions of Hecke irreducible subspaces inside the new subspace.

Endomorphism algebras. We can also see some interesting patterns in the distribution of endomorphism algebras of modular abelian surfaces over totally real fields. First, we have the following conjecture of Coleman (see e.g. Bruin–Flynn–González–Rotger [4, Conjecture $C(e,g)$].

Conjecture 4.1. *Let F be a totally real field and let $g \in \mathbb{Z}_{\geq 1}$. Then the set*

$$\{E = \mathrm{End}(X)_{\mathbb{Q}} : [E:\mathbb{Q}] = g \text{ and } X \text{ an abelian variety over } F \text{ with } \dim(X) = g\}$$

is finite.

In particular, there should be only finitely many real quadratic fields E that occur as the endomorphism algebra of a GL_2-type abelian surface ($g = 2$) over a fixed totally real field F. (Our data becomes a bit too sparse to say anything for $g \geq 3$ except for a few fields.) There are a total of 37444 Hilbert modular forms whose Hecke eigenvalue field has degree $g = 2$; corresponding to such a form f over F is an isogeny class of abelian surfaces X over F whose common endomorphism algebra $E = \mathrm{End}(X)_{\mathbb{Q}}$ is a quadratic field.

Hilbert modular forms over F whose Hecke field E is an imaginary quadratic field correspond to Hecke Grossencharacters. The largest absolute discriminant of an imaginary quadratic field observed was $E = \mathbb{Q}(\sqrt{-288})$, occurring for a Hilbert modular form over $F = \mathbb{Q}(\sqrt{440})$, a field with class number 2.

More interesting are the cases when the Hecke field E is a real quadratic field. The largest such real quadratic field observed as a Hecke field was $E = \mathbb{Q}(\sqrt{1260})$, occurring over the totally real quartic field F of discriminant 11025. The distribution of the quadratic Hecke fields for a fixed field F shows rapid decay by discriminant; the following histogram of the quadratic Hecke fields for $F = \mathbb{Q}(\sqrt{5})$ is typical.

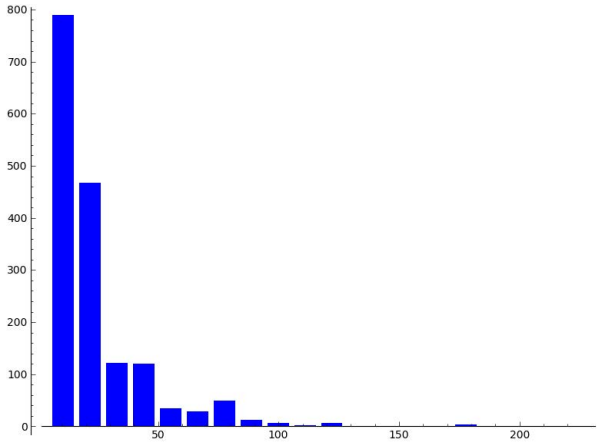

The largest discriminant field for $F = \mathbb{Q}(\sqrt{5})$ is $E = \mathbb{Q}(\sqrt{228})$, and it occurs just once with level norm 3420 (up to the computed level norm 4999). We tabulate in a similar way the following maximum discriminant quadratic Hecke fields for the first few F:

d_F	$\max(d_E)$
5	228
8	384
12	192
13	252
17	192
21	588
24	288
28	448

Put together, this data gives some weak evidence for Conjecture 4.1—at the very least, it shows that the large discriminant fields occur with very low density.

References

[1] B. J. Birch and W. Kuyk (eds.), *Modular functions of one variable IV*, Lecture Notes in Mathematics, vol. 476, Springer-Verlag, 1975.

[2] Jonathan Bober, Alyson Deines, Ariah Klages-Mundt, Benjamin LeVeque, R. Andrew Ohana, Ashwath Rabindranath, Paul Sharaba, and William Stein, *A database of elliptic curves over $\mathbb{Q}(\sqrt{5})$: a first report*, ANTS X: Proceedings of the Tenth Algorithmic Number Theory Symposium, eds. Everett W. Howe and Kiran S. Kedlaya, Open Book Series 1, 2013, 145–166.

[3] Wieb Bosma, John Cannon, and Catherine Playoust, *The Magma algebra system. I. The user language*, J. Symbolic Comput. **24** (1997), no. 3–4, 235–265.

[4] Nils Bruin, E. Victor Flynn, Josep González and Victor Rotger, *On finiteness conjectures for modular quaternion algebras*, Math. Proc. Cam. Phil. Soc. **141** (2006), 383–408.

[5] Henri Cohen, *A course in computational number theory*, Springer-Verlag, Berlin, 1993.

[6] J.E. Cremona, *Algorithms for modular elliptic curves*, 2nd ed., Cambridge, 1997.

[7] John Cremona, *The elliptic curve database for conductors to 130000*, Algorithmic number theory (Berlin, 2006), Lecture Notes in Comput. Sci., vol. 4076, Springer, Berlin, 2006, 11–29.

[8] L. Dembélé, *Quaternionic Manin symbols, Brandt matrices and Hilbert modular forms*, Math. Comp. **76** (2007), no. 258, 1039–1057.

[9] L. Dembélé and S. Donnelly, *Computing Hilbert modular forms over fields with nontrivial class group*, Algorithmic number theory (Banff, 2008), eds. Alfred van der Poorten and Andreas Stein, Lecture notes in Comput. Sci., vol. 5011, Springer, Berlin, 2008, 371–386.

[10] Lassina Dembélé and John Voight, *Explicit methods for Hilbert modular forms*, Elliptic curves, Hilbert modular forms and Galois deformations, Birkhauser, Basel, 2013, 135–198.

[11] E. Freitag, *Hilbert Modular Forms*, Springer-Verlag, Berlin, 1990.

[12] Matthew Greenberg and John Voight, *Computing systems of Hecke eigenvalues associated to Hilbert modular forms*, submitted.

[13] Markus Kirschmer and John Voight, *Algorithmic enumeration of ideal classes for quaternion orders*, SIAM J. Comput. (SICOMP) **39** (2010), no. 5, 1714-1747.

[14] The LMFDB Collaboration, *The L-functions and Modular Forms Database*, `http://www.lmfdb.org`, 2016.

[15] T. R. Shemanske and L. Walling, *Twists of Hilbert modular forms*, Trans. Amer. Math. Soc. **338** (1993), no. 1, 375–403.

[16] William Stein, *Modular forms database*, `http://modular.math.washington.edu/Tables`, 2004.

[17] William Stein, *Modular forms: a computational approach*, Grad. Studies in Math., vol. 79, American Mathematical Society, Providence, RI, 2007.

[18] John Voight, *Hilbert modular forms and modular elliptic curves*, `http://math.dartmouth.edu/ jvoight/hmf.html`.

[19] John Voight, *Computing fundamental domains for Fuchsian groups*, J. Théorie Nombres Bordeaux **21** (2009), no. 2, 467–489.

[20] John Voight, *Computing automorphic forms on Shimura curves over fields with arbitrary class number*, Algorithmic number theory (ANTS IX, Nancy, France, 2010), eds. Guillaume Hanrot, Francois Morain, and Emmanuel Thomé, Lecture Notes in Comp. Sci., vol. 6197, Springer, Berlin, 2010, 357–371.

School of Mathematics and Statistics F07, University of Sydney, Sydney, NSW 2006, Australia
E-mail address: `donnelly@maths.usyd.edu.au`

Department of Mathematics, Dartmouth College, 6188 Kemeny Hall, Hanover, NH 03755, USA
E-mail address: `jvoight@gmail.com`